

From Wide to Deep: Dimension Lifting Network for Parameter-efficient Knowledge Graph Embedding

Borui Cai, Yong Xiang, *Senior Member, IEEE*, Longxiang Gao, *Senior Member, IEEE*, Di Wu, *Member, IEEE*
He Zhang, Jiong Jin, *Member, IEEE* and Tom Luan, *Senior Member, IEEE*

Abstract—Knowledge graph embedding (KGE) that maps entities and relations into vector representations is essential for downstream applications. Conventional KGE methods require high-dimensional representations to learn the complex structure of knowledge graph, but lead to oversized model parameters. Recent advances reduce parameters by low-dimensional entity representations, while developing techniques (e.g., knowledge distillation or reinvented representation forms) to compensate for reduced dimension. However, such operations introduce complicated computations and model designs that may not benefit large knowledge graphs. To seek a simple strategy to improve the parameter efficiency of conventional KGE models, we take inspiration from that deeper neural networks require exponentially fewer parameters to achieve expressiveness comparable to wider networks for compositional structures. We view all entity representations as a single-layer embedding network, and conventional KGE methods that adopt high-dimensional entity representations equal widening the embedding network to gain expressiveness. To achieve parameter efficiency, we instead propose a deeper embedding network for entity representations, i.e., a narrow entity embedding layer plus a multi-layer dimension lifting network (LiftNet). Experiments on three public datasets show that by integrating LiftNet, four conventional KGE methods with 16-dimensional representations achieve comparable link prediction accuracy as original models that adopt 512-dimensional representations, saving 68.4% to 96.9% parameters.

Index Terms—Knowledge graph embedding, deep neural network, parameter-efficiency, representation learning.

I. INTRODUCTION

KNOWLEDGE graphs containing a large number of facts benefit various downstream applications, ranging from open-domain question answering [1], content-based recommender systems [2], to text-centric information retrieval [3]. A fact in a knowledge graph is represented as a triple, which includes a head/subject entity, a tail/object entity, and the relation identifying the relationship between them; for example, $\{Da\ Vinci, painted, Mona\ Lisa\}$. Due to the complex

This work has been submitted to the IEEE for possible publication. Copyright may be transferred without notice, after which this version may no longer be accessible.

B. Cai and Y. Xiang are with the School of Information Technology, Deakin University, VIC 3125, Australia. E-mail: {b.cai,yong.xiang}@deakin.edu.au.

L. Gao is with Qilu University of Technology (Shandong Academy of Sciences), China. E-mail: gaolx@sdas.org.

D. Wu is with School of Mathematics, Physics and Computing, University of Southern Queensland, QLD 4350, Australia. E-mail: di.wu@unisoq.edu.au.

H. Zhang is with CNPIEC KEXIN LTD, China. E-mail: zhanghe@kxsz.net.

J. Jin is with the School of Science, Computing and Engineering Technologies, Swinburne University of Technology, VIC 3122, Australia. E-mail: jiongjin@swin.edu.au.

T. Luan is with the School of Cyber Engineering, Xidian University, Xi'an, 710071, China. E-mail: tom.luan@xidian.edu.cn.




Fig. 1. In (a), conventional KGE models that use high-dimensional entity representations equal to enlarging the width of the embedding layer. But we tend to achieve parameter efficiency by increasing the depth of the embedding network, i.e., a narrower embedding layer (low-dimensional entity representations) plus the LiftNet as shown in (b).

multi-relational structure, directly applying knowledge graphs to downstream tasks is difficult. Knowledge Graph Embedding (KGE) methods [4] instead map entities and relations of a knowledge graph into vector representations, which preserve structural information of the knowledge graph and are convenient to be used by downstream tasks.

To accurately preserve the complex structural information of knowledge graphs, conventional KGE methods (e.g., TransE [4]) seek to increase the embedding dimension for better expressiveness and adopt high-dimensional entity/relation representations. However, since the scale of model parameters grows linearly with the representation dimension, this strategy leads to huge memory consumption, especially for large-scale real-life knowledge graphs. For example, when the embedding dimension is 512, a KGE model for Google knowledge graph [5] (covering 5 billion entities by 2020) requires over 2,560 billion model parameters; that is even 14 times larger than GPT-3 (one of the largest language models) [6]. This problem greatly limits KGE models on resource-constrained platforms.

Some methods attempt to improve parameter efficiency by reducing the embedding dimension, and the key problem is how to maintain model expressiveness. Meanwhile, since entities normally are much more than relations in many knowledge graphs, they mostly reduce the dimension of entity representations. One type of method adopts knowledge distillation to transfer knowledge from pre-trained high-dimensional KGE models to low-dimensional entity representations [7], [8]. However, pre-trained high-dimensional KGE models are not always available, and such a training process also demands a huge memory footprint. Other methods reinvent more expressive representation forms, e.g., hyperbolic representation

[9], or compositional anchor nodes [10]; but these operations introduce complicated geometric computations and model designs that may not benefit large knowledge graphs.

In search of a simple way to reduce entity parameters and hence improve the parameter efficiency of conventional KGE models, we take inspiration from earlier discussions on wider vs deeper neural networks. Study shows that deeper neural networks require exponentially fewer model parameters than wider networks to provide similar expressiveness for compositional structures [11], [12], which widely exists in knowledge graphs [13]. From this perspective, we view the concatenation of all entity representations as a single-layer embedding network. Then, a conventional KGE model that requires high-dimensional entity representations equals a rather wide network, i.e., by enlarging the width of the embedding layer to gain model expressiveness (Fig. 1 (a)). Instead, we argue to simply achieve parameter efficiency by increasing the depth of the embedding network (Fig. 1 (b)). Specifically, our deeper entity embedding network includes 1) an input layer that receives low-dimensional (\hat{n}) entity representations; 2) multiple hidden layers following the input layer to increase the model expressiveness; 3) an output layer that produces high-dimensional (n) entity representations, which can be directly adopted by existing KGE methods. We denote the entity embedding network as dimension-lifting network (LiftNet) since it lifts low-dimensional entity representations to high-dimensional. We refrain applying LiftNet to relation to accommodate conventional KGE models that design relation as other type of operations, e.g., translation on hyperplanes [14]. The number of parameters for LiftNet-based entity representations becomes $|\mathcal{E}|\hat{n} + P_{net}$. $|\mathcal{E}|$ is the number of entities in the knowledge graph and P_{net} is the model parameter size of LiftNet, which is negligible compared to $|\mathcal{E}|\hat{n}$ for large knowledge graphs. LiftNet can be conveniently integrated with conventional KGE models to make them parameter-efficient. We evaluate the effectiveness by integrating LiftNet with four strong KGE methods and running experiments on three knowledge graph datasets of different sizes. The results show that by integrating with LiftNet, conventional KGE methods only require 16-dimensional entity representations to achieve link prediction accuracy comparable to original models of 512-dimensional, saving 68.4% to 96.9% model parameters.

II. RELATED WORK

A. Knowledge graph embedding

KGE methods learn vector representations for knowledge graphs, and we roughly categorize them into three types. First, distance-based methods describe a fact with mathematical operations, e.g., TransE defines a relation as the translation. To better model N-N relations, TransH [14] and STTransE [15] project entities to relation-aware subspace with hyperplanes and matrices, respectively. Operations in the complex space [16] or the polar coordinate system [17] are also introduced to improve flexibility. Second, tensor factorization methods, e.g., RESCAL [18], model the knowledge graph as a three-way tensor, and apply tensor decomposition to obtain entity/relation representations. Later methods further improve the

effectiveness of decomposition. For example, DistMult [19] applies an efficient diagonal relation matrix, ComplEx [20] introduces complex embedding to capture relation asymmetry, and Simple [21] addresses the independence of entity embedding. Third, deep learning methods adopt deep neural networks to capture the complex relationships of entities and relations. ConvE [22], and CapsE [23] learn the complex interactions between entities and relations through convolutional layers and capsule layers, respectively. CompGCN [24] leverages GCN layers with entity-relation composition operations to capture interactions among entities and relations.

B. Parameter-efficient knowledge graph embedding

Since the scale of model parameters is linear with the embedding dimensions, many works focus on employing different techniques to improve the effectiveness of low-dimensional entity representations. Knowledge distillation is adopted to train low-dimensional representations with multiple pre-trained KGE models [7], which shows better performance than training from scratch. The number of pre-trained models is further reduced to one [8] by introducing the dual influence between the pre-trained and the target models. Meanwhile, to accelerate the distillation, the low-dimensional model training is replaced by more efficient feature pruning of pre-trained models [25]. However, the obtained low-dimensional entity representations still show degraded performance. To improve the performance, ROTH [9] adopts more flexible hyperbolic space (compared to Euclidean space) for low-dimensional entity representations, but that also brings more complex hyperbolic geometry [26]. To avoid that, contrastive learning is employed to flexibly control the strength of penalties for easy/difficult instances [27]. Other works reinvent the entity representations as more expressive forms. For example, through learning a vocabulary of abstractive codewords that capture shared features of entities, entity representations become the composition of relevant codewords [10], [13].

III. THE PROPOSED METHOD

A. Preliminaries

A knowledge graph is a multi-relational graph denoted as $\mathcal{G} = (\mathcal{E}, \mathcal{R}, \mathcal{D})$. \mathcal{E} and \mathcal{R} are the collections of entities and relations, respectively, and \mathcal{D} is the collection of facts contained in the knowledge graph. A fact $d \in \mathcal{D}$ is a triple denoted as $\{h, r, t\}$. In d , $h \in \mathcal{E}$ is the head entity, $t \in \mathcal{E}$ is the tail entity, and $r \in \mathcal{R}$ is the relation between the head entity and tail entity. Knowledge graph embedding aims at learning vector representations for entities and relations while preserving the structural information of the knowledge graph. We denote the vector representations corresponding to $\{h, r, t\}$ as $\{e_h, r_r, e_t\}$. Conventional KGE methods map h, r , and t into the same latent space, i.e., $e_h, r_r, e_t \in \mathbb{R}^n$, and learn entity/relation representations by minimizing score functions defined as geometry operations; for example, $l(h, r, t) = \|e_h + r_r - e_t\|_p$ (TransE).

Most KGE methods require relatively high embedding dimensions to sufficiently express the structural information of knowledge graph [9]; however, that leads to oversized model

parameters. Adopting low-dimensional ($\hat{e} \in \mathbb{R}^{\hat{n}}, \hat{n} \ll n$) entity representations can significantly reduce model parameters but results in low performance due to decreased expressiveness. To achieve parameter-efficient KGE, we aim at using low-dimensional representations ($\hat{e} \in \mathbb{R}^{\hat{n}}, \hat{n} \ll n$), without degrading the performance. Following previous methods [13], [28], we focus on reducing the dimensions of entity representations since the number of entities is much larger than that of relations in many real-life knowledge graphs.

B. Dimension Lifting Network

As discussed above, we see the concatenation of entity representations as an embedding layer. To achieve parameter efficiency, we aim at replacing the wide embedding layer of high-dimensional entity representations ($\{e_1 || e_2 || \dots || e_{|\mathcal{E}|}\}$) with a deeper embedding network; that is, a narrow embedding layer of low-dimensional entity representations ($\{\hat{e}_1 || \hat{e}_2 || \dots || \hat{e}_{|\mathcal{E}|}\}$) and a multi-layer dimension lifting network (LiftNet). $\{\hat{e}_1 || \hat{e}_2 || \dots || \hat{e}_{|\mathcal{E}|}\}$ and LiftNet are fully-connected; that means LiftNet can be interpreted as a dimension lifting function $f(*)$ that lifts $\hat{e}_i \in \mathbb{R}^{\hat{n}}$ to $e = f(\hat{e}_i) \in \mathbb{R}^n$. $f(\hat{e}_i)$ is used together with the original relation representations for score measurements and inference.




Fig. 2. The structure of LiftNet. \hat{e} is the low-dimensional input entity representation, and LiftNet uses n TC layers to progressively lift it to high-dimensional output e .

The main task is to design an effective $f(*)$ for KGE. An intuitive choice of $f(*)$ is multiple fully connected (FC) layers; however, FC layers require large numbers of parameters and are prone to overfitting for KGE [22]. Inspired by image processing, we refer to feature upsampling techniques for dimension lifting. Specifically, we adopt transposed convolution (TC) layer in LiftNet. A TC layer broadcasts the input elements via kernels, thus increasing the dimension of the output. Different from traditional upsampling methods (e.g., nearest-neighbor, bilinear, and bicubic interpolation [29]), TC can capture the interactions among the parameters of entity representations. The structure of LiftNet is shown in Fig. 2.

LiftNet adopts multiple TC layers to progressively lift the low-dimensional \hat{e} to the high-dimensional e (we find in experiments that LiftNet with only two FC layers already shows promising performance on three public knowledge graph datasets). LiftNet reshapes \hat{e} as square feature maps in the forward pass and feeds them into the TC layers. After that, the representation is lifted into $\mathbb{R}^{c \times m_1 \times m_2}$, where c is the output channel number and $m_1 \times m_2$ is the size of output feature map. The output of the last TC layer is then reshaped

back to the expected high-dimensional e . We use Tanh as the non-linear activation function between TC layers and also the output layer, considering that it is zero-centered and can capture complex feature patterns. We choose Adam [30] as the optimizer for the model training.

LiftNet is model-agnostic and can conveniently integrate with many conventional KGE methods for parameter-efficient representation learning, e.g., TransE, DistMult, by simply replacing the entity representation e with $f(\hat{e})$. The score function of LiftNet-based method becomes $l(h, r, t) = l(f(\hat{e}_h), r_r, f(\hat{e}_t))$, with $f(\hat{e}_h), r_r, f(\hat{e}_t) \in \mathbb{R}^n$. The parameter size of entity/relation representation for LiftNet-based methods is $|\mathcal{E}|\hat{n} + |\mathcal{R}|n + P_{net}$, where P_{net} is the size of LiftNet parameters (mainly the transposed convolution kernels in TC layers) and $P_{net} \ll |\mathcal{E}|\hat{n}$. Therefore, for many real-world KG datasets that have far more entities than relations ($|\mathcal{E}| \gg |\mathcal{R}|$), LiftNet-based KGE methods only require approximately $\frac{\hat{n}}{n}$ parameters of original models that need high-dimensional entity representations.

TABLE I
STATISTICS OF THE DATASETS.

Dataset	#E	#R	Train/Valid/Test
UMLS	135	46	5,327/569/633
WN18RR	40,943	11	86,835/3,034/3,134
FB15K237	14,541	237	272,115/17,535/20,466

IV. EXPERIMENTS

A. Experiment Setup

a) *Datasets*: We adopt three public knowledge graph datasets widely used for KGE and link prediction, i.e., UMLS [31], WN18RR [4], and FB15K237 [32]. Detailed statistics of these datasets are summarized in Table I. UMLS is a small dataset derived from the repository of biomedical vocabularies. WN18RR is a subset of WordNet, which describes hyponym and hypernym relations among words, and it removes inverse relations to avoid test leakage. FB15K237 adopts a similar process to remove nearly identical and inverse relations in FB15K [4], which contains entity-relation triples of Freebase.

b) *Evaluation metrics*: We evaluate the performance of KGE with link prediction [4], which is a popular application of KGE. Specifically, the task is to predict missing entity given a query ($\{h, r, ?\}$ or $\{?, r, t\}$). Following existing works [19], [20], we use mean reciprocal ranking (MRR) and H@k to measure the prediction accuracy. MRR is the average inverse rank for all test triples, and H@k is the percentage of ranks lower than or equal to k. The maximum values of MRR and H@k are both 1, and the higher MRR or H@k, the better the performance. We adopt the filtered setting [9] to exclude candidate triples that have been seen in training, validation, and testing sets.

c) *Conventional KGE models*: We choose a set of strong conventional KGE models to show the effectiveness of the proposed LiftNet method. That includes TransE, TransH, DistMult, and ComplEx. TransE and TransH are translational models that adopt distance measurements for related entities

TABLE II

LINK PREDICTION ACCURACY OF LIFTNET METHODS AND CONVENTIONAL KGE METHODS (THE BEST IN BOLD AND THE SECOND-BEST UNDERLINED).

Methods (dim)	UMLS			WN18RR			FB15K237		
	MRR	H@1	H@10	MRR	H@1	H@10	MRR	H@1	H@10
TransE (16)	.568	.372	.844	.102	.020	.281	.138	.074	.261
TransE (512)	<u>.661</u>	<u>.423</u>	.986	.172	.097	<u>.392</u>	.208	<u>.076</u>	.445
LN-TransE (16)	.782	.631	<u>.970</u>	<u>.164</u>	<u>.082</u>	.406	<u>.205</u>	.094	<u>.412</u>
TransH (16)	.668	.379	.897	.100	.014	.275	.146	.072	.261
TransH (512)	<u>.674</u>	<u>.513</u>	<u>.988</u>	<u>.167</u>	<u>.030</u>	<u>.393</u>	<u>.208</u>	<u>.084</u>	.445
LN-TransH (16)	.848	.730	.989	.174	.032	.417	.235	.134	<u>.429</u>
DistMult (16)	.100	.012	.270	.015	.004	.034	.031	.011	.057
DistMult (512)	<u>.655</u>	<u>.558</u>	<u>.846</u>	.388	.329	.477	<u>.234</u>	<u>.146</u>	.407
LN-DistMult (16)	.710	.614	.891	<u>.246</u>	<u>.296</u>	<u>.427</u>	.245	.166	<u>.405</u>
ComplEx (16)	.751	.608	.968	.112	.047	.235	.094	.026	.237
ComplEx (512)	.870	.781	<u>.977</u>	.409	.369	.471	<u>.259</u>	<u>.171</u>	.437
LN-ComplEx (16)	<u>.806</u>	<u>.684</u>	.981	<u>.332</u>	<u>.262</u>	<u>.441</u>	.266	.178	<u>.432</u>

TABLE III

LINK PREDICTION ACCURACY OF PARAMETER-EFFICIENT KGE METHODS (THE BEST IN BOLD AND THE SECOND-BEST UNDERLINED).

Methods	UMLS		WN18RR		FB15K237	
	MRR	H@10	MRR	H@10	MRR	H@10
DualDE	.763	.967	.168	.389	.157	.374
LightKG	.822	<u>.970</u>	.175	.401	.233	.418
GreenKGC	.727	.959	<u>.250</u>	<u>.429</u>	<u>.252</u>	.463
LN-ComplEx	<u>.806</u>	.981	.332	.441	.266	<u>.432</u>

and their relations. DistMult and ComplEx aim at semantic matching and adopt tensor decomposition in real and complex spaces, respectively. Their scoring functions are summarized in Table IV. By replacing e with $f(\hat{e})$ in the above four methods, we implement LN-TransE, LN-TransH, LN-DistMult, and LN-ComplEx for parameter-efficient KGE.

TABLE IV

SCORING FUNCTIONS OF EXPERIMENTED KGE METHODS.

Methods	Scoring Function $l(h, r, t)$
TransE	$\ e_h + r_r - e_t\ _p$
TransH	$\ e_{h,\perp} + r_r - e_{t,\perp}\ _p, e_{\perp} = e - w_r e w_r^T$
DistMult	$\langle e_h, r_r, e_t \rangle$
ComplEx	$\langle e_h, r_r, e_t \rangle, r_r \in \mathcal{C}^n$

d) Implementation details: We implement all the methods with OpenKE [33], which is a pytorch-based open-source framework for knowledge embedding¹. We run TransE, TransH, DistMult, and ComplEx with low-dimensional (16) and high-dimensional (512) embedding dimensions to show the difference in link prediction performance. Meanwhile, we implement a two-layer LiftNet for LiftNet-based methods, which lifts the dimension of entity representations from 16 to 512. In LiftNet, the parameters of the two TC layers as set as {InChan:1, OutChan:4, Kernel:3} and {InChan:4, OutChan:8, Kernel:3}, respectively, with stride=1 and padding=0. We fix the random seed for all experiments and use MRR of the

validation set to find the optimal learning rate from {0.01, 0.05, 0.1, 0.5}. The maximum epoch is 500.

B. Main Results

a) Link prediction accuracy: The accuracy of link prediction of the proposed LifeNet-based methods and compared conventional methods are shown in Table II. We observe that TransE, TransH, DistMult, and ComplEx all obtain higher link prediction accuracy with 512 embedding dimensions than with 16 embedding dimensions, due to the increased expressiveness of entity/relation representations. Especially on WN18RR and FB15K237 datasets, the accuracy of 512-dimensional models is far better than 16-dimensional models. Meanwhile, LN-TransE, LN-TransH, LN-DistMult, and LN-ComplEx, which also adopt 16-dimensional entity representations, achieve more accurate prediction results than corresponding methods that adopt 16-dimensional representations. In fact, their results are comparable to original models with 512-dimensional representations. We observe that on UMLS and FB15K237 datasets, MRR, H@1, and H@10 obtained by LN-TransH are even higher than those of original TransH with 512 embedding dimension; LN-DistMult also outperforms DistMult with 512 embedding dimension on UMLS dataset. For the rest results, LiftNet-based methods only show link prediction accuracy that are slightly lower than original methods with 512 embedding dimension. We also observe that LiftNet methods perform better on UMLS and FB15K237 datasets than on WN18RR, because UMLS and FB15K237 contain richer structural information, i.e., the average edges per entity (39.5 and 18.7) are much higher than that of WN18RR (2.1).

We further compare the performance of LiftNet with existing parameter-efficient KGE methods on the three datasets. We select three strong parameter-efficient KGE methods of different types, i.e., DualDE [7], LightKG [34], and GreenKGC [25], which performs distillation, quantization, and pruning, respectively. For a fair comparison, we adopt the same random negative sampling [35] and unify the base model of all parameter-efficient methods as ComplEx (since ComplEx performs better than TransE, TransH, and DistMult on the three datasets). We follow the model setup of

¹Codes are available at <https://github.com/brcai/LiftNet>

LightKE [34] and adopt 16-dimensional entity representations for the rest parameter-efficient models. As shown in Table III, LN-ComplEx generally achieves more accurate link prediction than DualDE, LightKG, and GreenKGC. LN-ComplEx achieves the best link prediction accuracy on WN18RR dataset (both MRR and H@10), and the best or the second best on UMLS and FB15K237 datasets. In particular, LN-ComplEx outperforms DualDE on all datasets, and is slight lower than LightKG and GreenKGC on UMLS (MRR is smaller for 0.016) and FB15K237 (H@10 is smaller for 0.031), respectively.

TABLE V
PARAMETER EFFICIENCY (M IS SHORT FOR MILLION).

	UMLS	WN18RR	FB15K237
TransE (512)	0.092M	20.968M	7.566M
LN-TransE (16)	0.026M	0.661M	0.354M
<i>Percentage</i>	28.2%	3.1%	4.7%
TransH (512)	0.116M	20.974M	7.687M
LN-TransH (16)	0.049M	0.667M	0.475M
<i>Percentage</i>	42.2%	3.2%	6.2%
DistMult (512)	0.092M	20.968M	7.566M
LN-DistMult (16)	0.026M	0.661M	0.354M
<i>Percentage</i>	28.2%	3.1%	4.7%
ComplEx (512)	0.185M	41.926M	15.132M
LN-ComplEx (16)	0.051M	1.321M	0.708M
<i>Percentage</i>	27.5%	3.1%	4.7%

TABLE VI
MEMORY USAGE (MB IS SHORT FOR MEGABYTE).

	UMLS	WN18RR	FB15K237
TransE (512)	0.354MB	80.022MB	28.864MB
LN-TransE (16)	0.101MB	2.523MB	1.353MB
<i>Percentage</i>	28.5%	3.1%	4.7%
TransH (512)	0.445MB	80.043MB	29.327MB
LN-TransH (16)	0.191MB	2.544MB	1.816MB
<i>Percentage</i>	42.9%	3.2%	6.2%
DistMult (512)	0.354MB	80.022MB	28.864MB
LN-DistMult (16)	0.101MB	2.523MB	1.353MB
<i>Percentage</i>	28.5%	3.1%	4.7%
ComplEx (512)	0.708MB	160.043MB	57.728MB
LN-ComplEx (16)	0.199MB	5.043MB	2.704MB
<i>Percentage</i>	28.1%	3.1%	4.7%

b) Parameter efficiency: In Table V, we show the parameter efficiency of LiftNet-based methods on the three datasets. The results are shown pair-wisely, i.e., the numbers of parameters required by 512-dimensional KGE models and those of corresponding LiftNet models, because they achieve similar link prediction accuracy. KGE methods with 16-dimensional representations are excluded due to their unsatisfying link prediction accuracy. As the results show, the LiftNet-based methods save 68.4% to 96.9% model parameters or the original models. Since the parameter size of LiftNet is negligible (324 parameters), the level of parameter efficiency is mainly affected by the ratio of entities to relations in the dataset. On the smallest UMLS dataset, which has the smallest entity/relation ratio (around 3), LiftNet-based methods require approximately

28% ~ 42% model parameters of original KGE models. However, on larger WN18RR and FB15K237 datasets that have much larger entity/relation ratios (around 3722 and 61, respectively), they only require approximately 3.1% and 5.1% model parameters of the original KGE methods to achieve comparable link prediction accuracy. This makes LiftNet promising for real-world knowledge graphs such as Google knowledge graph that has billions of entities. We compare the memory usage of 512-dimensional KGE models and the corresponding LiftNet models. The results in Table VI show that adopting LiftNet can significantly reduce the memory usage, e.g., from around 160 MB (512-dimensional ComplEx) to 5 MB (16-dimensional LN-ComplEx) on WN18RR dataset. Meanwhile, the percentage of memory reduction is consistent to that of model parameters (Table V), especially on larger WN18RR and FB15K237 datasets, proving the reduction of memory usage is due to the fewer model parameters of LiftNet.

We further discuss the impact of LiftNet on training complexity since it introduces additional neural layers to obtain entity representations. Specifically, we train each method for 100 epochs and then provide the average training time taken per epoch. As shown in Table VII, by adopting LiftNet, the ratio of training time increases ranges from 1.08 to 1.33 times, with respect to 0.878 seconds to 0.814 seconds (LN-TransE to TransE, WN18RR dataset) and 0.052 seconds to 0.039 seconds (LN-DistMult to DistMult, UMLS dataset). The above results show that training time increases is insignificant compared to the improved parameter efficiency.

TABLE VII
AVERAGE TRAINING TIME FOR ONE EPOCH (SECONDS).

	UMLS	WN18RR	FB15K237
TransE (512)	0.032	0.814	1.971
LN-TransE (16)	0.036	0.878	2.107
<i>Ratio</i>	1.13	1.08	1.07
TransH (512)	0.064	1.073	3.368
LN-TransH (16)	0.081	1.289	4.156
<i>Ratio</i>	1.27	1.20	1.23
DistMult (512)	0.039	0.783	2.392
LN-DistMult (16)	0.052	0.962	3.038
<i>Ratio</i>	1.33	1.23	1.27
ComplEx (512)	0.078	1.589	4.424
LN-ComplEx (16)	0.101	1.859	5.470
<i>Ratio</i>	1.29	1.17	1.24

C. Sensitivity Analysis

a) Input dimensions: The input dimensions highly affect the parameter efficiency of LiftNet-based methods; thus we analyze its influence on the model performance. We evaluate LN-TransE, LN-TransH, LN-DistMult, and LiftNet-ComplEx with four input dimensions {4, 16, 64, 256} on WN18RR dataset. To do that, we adjust the setups of TC layers (channel number, kernel size, and stride) accordingly to lift the dimensions of entity representations to 512. Due to the characteristics of TC layers, a higher input dimension requires a lower number of parameters (fewer kernels) in LiftNet.

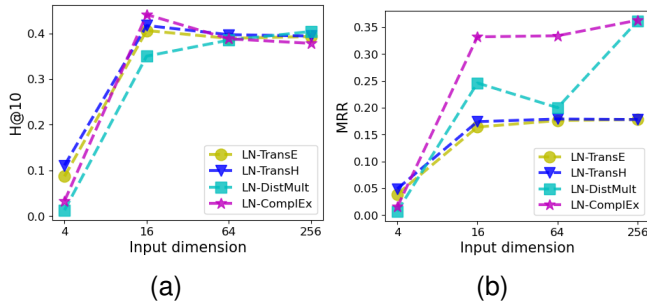


Fig. 3. Link prediction accuracy w.r.t. input dimensions.

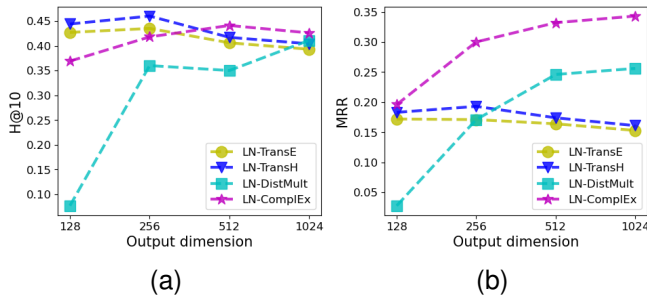


Fig. 4. Link prediction accuracy w.r.t. output dimensions.

The results of LiftNet-based methods for knowledge graph link prediction (accuracy measured by H@10 and MRR) are shown in Fig. 3. Generally, on WN18RR datasets, we observe the link prediction accuracy increases with higher input dimension, and the increase is significant from 4-dimension to 16-dimension. However, after that, the accuracy only slightly increases or even drops, e.g., LN-CompLex and LN-DistMult in Fig. 3 (a), with higher input dimensions partly due to over-parameterization. That shows the effectiveness of LiftNet with low-dimensional entity representations.

b) Output dimension: We show the sensitivity of the performance of LiftNet-based methods regarding the output dimensions on WN18RR dataset. To do that, we vary the output dimensions from 128 to 1024 and adjust the setups of the two TC layers in LiftNet accordingly. The results in Fig. 4 show two different trends. First, LN-TransE and LN-TransH show increased H@10 in Fig. 4 (a) and MRR in Fig. 4 (b) with higher output dimensions and then slightly decrease; while the results of LN-DistMult and LN-CompLex mostly increase. The difference is caused by their different scoring functions that lead to different sensitivity to over-parameterization; additive scoring functions for LN-TransE and LN-TransH and multiplicative scoring functions for LN-DistMult and LN-CompLex. The above results generally support why LiftNet needs to “lift” output dimensions, rather than maintain or even reduce them.

c) Number of layers: We further analyze the sensitivity of LiftNet regarding the number of TC layers on WN18RR dataset. For demonstration, we vary the number of TC layers in LiftNet-based methods from one to four. Note one layer LiftNet is still different from the linear transformation used in [28] as it has non-linear activation on the output layer. We adjust the setups of TC layers by increasing/decreasing the

kernel size accordingly to ensure the output dimensions are always 512. The results with respect to link prediction accuracy (H@10) are shown in Table VIII. We see that the results of LN-TransE and LN-TransH are relatively stable, introducing more TC layers only slightly increases model performance. While LN-DistMult and LN-CompLex are more sensitive to different layer numbers, their performance collapses with 4 TC layers partly due to training difficulties and overfitting.

TABLE VIII
THE LINK PREDICTION ACCURACY (BY MRR) OF LIFTNET-BASED METHODS W.R.T. THE NUMBER OF TC LAYERS.

Layers	LN-TransE	LN-TransH	LN-DistMult	LN-CompLex
1	.395	.412	.088	.474
2	.406	.417	.427	.441
3	.406	.422	.358	.375
4	.411	.421	.004	.010

TABLE IX
THE LINK PREDICTION ACCURACY (BY MRR) OF LIFTNET VARIANTS THAT ADOPT FC LAYERS.

Layers	LN-TransE	LN-TransH	LN-DistMult	LN-CompLex
TC (2)	.412	.429	.405	.432
FC (2)	.352	.402	.408	.411
FC (3)	.353	.405	.260	.289
FC (4)	.355	.404	.124	.180

D. LiftNet Variants

In LiftNet, we adopt TC layers to progressively lift the dimensions. To demonstrate the effectiveness of such a design, we implement LiftNet variants with fully connected (FC) layers for comparison. The experiment is conducted on the largest FB15K237 dataset, with accuracy measured by MRR. Specifically, we include LiftNet variants of 2 to 4 FC layers, and the results are shown in Table IX. We see that in most cases, LiftNet with 2 TC layers achieves more accurate link prediction results than LiftNet with 2 or more FC layers. We do not claim that TC layers are better than FC layers, but in KGE where the model is sensitive to insufficient expressiveness or over-parameterization, LiftNet with TC layers is relatively easier to achieve more accurate link prediction results.

V. CONCLUSION

In this paper, we propose a simple LiftNet method that helps improve the parameter efficiency of conventional KGE models. LiftNet adopts a multi-layer neural network to enhance the expressiveness of low-dimensional entity representations. Experiments conducted on three public knowledge graph datasets show that with LiftNet, conventional KGE models TransE, TransH, DistMult, and CompLex can save up to 68.4%, 96.9%, and 94.9% model parameters on UMLS, WN18RR, and FB15K237 datasets, respectively, in link prediction. One potential drawback of LiftNet is the increased training/inference time, due to the multi-layer back/forward propagation. However, experiments show that LiftNet requires

only two TC layers for the tested datasets, hence the additional run-time complexity pales in comparison to the improved parameter efficiency. Meanwhile, LiftNet still requires input entity representations to generate desired output, in future work, we will explore semantic encoders that generate entity representations from their textual information (e.g., names and descriptions) for higher model effectiveness.

REFERENCES

- [1] Q. Lang, X. Liu, and W. Jia, "Afs graph: Multidimensional axiomatic fuzzy set knowledge graph for open-domain question answering," *IEEE Transactions on Neural Networks and Learning Systems*, 2022.
- [2] Y. Du, X. Zhu, L. Chen, Z. Fang, and Y. Gao, "Metakg: Meta-learning on knowledge graph for cold-start recommendation," *IEEE Transactions on Knowledge and Data Engineering*, 2022.
- [3] Z. Yang, "Biomedical information retrieval incorporating knowledge graph for explainable precision medicine," in *Proceedings of the 43rd International ACM SIGIR Conference on Research and Development in Information Retrieval*, 2020, pp. 2486–2486.
- [4] A. Bordes, N. Usunier, A. Garcia-Duran, J. Weston, and O. Yakhnenko, "Translating embeddings for modeling multi-relational data," *Advances in Neural Information Processing Systems*, vol. 26, 2013.
- [5] D. Sullivan, "A reintroduction to our knowledge graph and knowledge panels," *The Keyword, Google, May*, vol. 20, 2020.
- [6] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell *et al.*, "Language models are few-shot learners," *Advances in Neural Information Processing Systems*, vol. 33, pp. 1877–1901, 2020.
- [7] Y. Zhu, W. Zhang, M. Chen, H. Chen, X. Cheng, W. Zhang, and H. Chen, "Dualde: Dually distilling knowledge graph embedding for faster and cheaper reasoning," in *Proceedings of the Fifteenth ACM International Conference on Web Search and Data Mining*, 2022, pp. 1516–1524.
- [8] K. Wang, Y. Liu, Q. Ma, and Q. Z. Sheng, "Mulde: Multi-teacher knowledge distillation for low-dimensional knowledge graph embeddings," in *Proceedings of the Web Conference 2021*, 2021, pp. 1716–1726.
- [9] I. Chami, A. Wolf, D.-C. Juan, F. Sala, S. Ravi, and C. Ré, "Low-dimensional hyperbolic knowledge graph embeddings," in *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, 2020, pp. 6901–6914.
- [10] M. Galkin, J. Wu, E. Denis, and W. L. Hamilton, "Nodepiece: Compositional and parameter-efficient representations of large knowledge graphs," *arXiv preprint arXiv:2106.12144*, 2021.
- [11] H. Mhaskar, Q. Liao, and T. Poggio, "When and why are deep networks better than shallow ones?" in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 31, no. 1, 2017.
- [12] H. W. Lin, M. Tegmark, and D. Rolnick, "Why does deep and cheap learning work so well?" *Journal of Statistical Physics*, vol. 168, pp. 1223–1247, 2017.
- [13] H. Wang, Y. Wang, D. Lian, and J. Gao, "A lightweight knowledge graph embedding framework for efficient inference and storage," in *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 2021, pp. 1909–1918.
- [14] Z. Wang, J. Zhang, J. Feng, and Z. Chen, "Knowledge graph embedding by translating on hyperplanes," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 28, no. 1, 2014.
- [15] D. Q. Nguyen, K. Sirts, L. Qu, and M. Johnson, "STransE: a novel embedding model of entities and relationships in knowledge bases," in *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, Jun. 2016, pp. 460–466.
- [16] Z. Sun, Z.-H. Deng, J.-Y. Nie, and J. Tang, "Rotate: Knowledge graph embedding by relational rotation in complex space," in *International Conference on Learning Representations*, 2019.
- [17] Z. Zhang, J. Cai, Y. Zhang, and J. Wang, "Learning hierarchy-aware knowledge graph embeddings for link prediction," in *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020, pp. 3065–3072.
- [18] Y. Lin, Z. Liu, M. Sun, Y. Liu, and X. Zhu, "Learning entity and relation embeddings for knowledge graph completion," in *Twenty-ninth AAAI Conference on Artificial Intelligence*, 2015.
- [19] B. Yang, W.-t. Yih, X. He, J. Gao, and L. Deng, "Embedding entities and relations for learning and inference in knowledge bases," *arXiv preprint arXiv:1412.6575*, 2014.
- [20] T. Trouillon, J. Welbl, S. Riedel, É. Gaussier, and G. Bouchard, "Complex embeddings for simple link prediction," in *International Conference on Machine Learning*. PMLR, 2016, pp. 2071–2080.
- [21] S. M. Kazemi and D. Poole, "Simple embedding for link prediction in knowledge graphs," *Advances in Neural Information Processing Systems*, vol. 31, 2018.
- [22] T. Dettmers, P. Minervini, P. Stenetorp, and S. Riedel, "Convolutional 2d knowledge graph embeddings," in *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 32, no. 1, 2018.
- [23] T. Vu, T. D. Nguyen, D. Q. Nguyen, D. Phung *et al.*, "A capsule network-based embedding model for knowledge graph completion and search personalization," in *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2019, pp. 2180–2189.
- [24] S. Vashishth, S. Sanyal, V. Nitin, and P. Talukdar, "Composition-based multi-relational graph convolutional networks," in *International Conference on Learning Representations*, 2020.
- [25] Y. C. Wang, X. Ge, B. Wang, and C.-C. J. Kuo, "GreenKGC: A lightweight knowledge graph completion method," in *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, 2023, pp. 10 596–10 613.
- [26] K. Wang, Y. Liu, D. Lin, and Q. Z. Sheng, "Hyperbolic geometry is not necessary: Lightweight euclidean-based models for low-dimensional knowledge graph embeddings," *arXiv preprint arXiv:2103.14930*, 2021.
- [27] K. Wang, Y. Liu, and Q. Z. Sheng, "Swift and sure: Hardness-aware contrastive learning for low-dimensional knowledge graph embeddings," in *Proceedings of the ACM Web Conference 2022*, 2022, pp. 838–849.
- [28] L. Chao, T. Wang, and W. Chu, "Pie: a parameter and inference efficient solution for large scale knowledge graph embedding reasoning," *arXiv preprint arXiv:2204.13957*, 2022.
- [29] H. Tang, L. Shao, P. H. Torr, and N. Sebe, "Local and global gans with semantic-aware upsampling for image generation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 768–784, 2022.
- [30] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [31] O. Bodenreider, "The unified medical language system (umls): integrating biomedical terminology," *Nucleic Acids Research*, vol. 32, no. suppl_1, pp. D267–D270, 2004.
- [32] K. Toutanova and D. Chen, "Observed versus latent features for knowledge base and text inference," in *Proceedings of the 3rd Workshop on Continuous Vector Space Models and Their Compositionality*, 2015.
- [33] X. Han, S. Cao, X. Lv, Y. Lin, Z. Liu, M. Sun, and J. Li, "Openke: An open toolkit for knowledge embedding," in *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, 2018, pp. 139–144.
- [34] H. Wang, Y. Wang, D. Lian, and J. Gao, "A lightweight knowledge graph embedding framework for efficient inference and storage," in *Proceedings of the 30th ACM International Conference on Information & Knowledge Management*, 2021, pp. 1909–1918.
- [35] T. Madushanka and R. Ichise, "Negative sampling in knowledge graph representation learning: A review," *arXiv preprint arXiv:2402.19195*, 2024.